

On The Power of Tree Projections: Structural Tractability of Enumerating CSP Solutions

Gianluigi Greco¹ and Francesco Scarcello²

Dept. of Mathematics¹ and DEIS², University of Calabria, 87036, Rende, Italy
ggreco@mat.unical.it, scarcello@deis.unical.it

Abstract. The problem of deciding whether CSP instances admit solutions has been deeply studied in the literature, and several structural tractability results have been derived so far. However, constraint satisfaction comes in practice as a computation problem where the focus is either on finding one solution, or on enumerating all solutions, possibly projected to some given set of output variables. The paper investigates the structural tractability of the problem of enumerating (possibly projected) solutions, where tractability means here computable with polynomial delay (WPD), since in general exponentially many solutions may be computed. A general framework based on the notion of tree projection of hypergraphs is considered, which generalizes all known decomposition methods. Tractability results have been obtained both for classes of structures where output variables are part of their specification, and for classes of structures where computability WPD must be ensured for any possible set of output variables. These results are shown to be tight, by exhibiting dichotomies for classes of structures having bounded arity and where the tree decomposition method is considered.

1 Introduction

1.1 Constraint Satisfaction and Decomposition Methods

Constraint satisfaction is often formalized as a homomorphism problem that takes as input two finite relational structures \mathbb{A} (modeling variables and scopes of the constraints) and \mathbb{B} (modeling the relations associated with constraints), and asks whether there is a homomorphism from \mathbb{A} to \mathbb{B} . Since the general problem is NP-hard, many restrictions have been considered in the literature, where the given structures have to satisfy additional conditions. In this paper, we are interested in restrictions imposed on the (usually said) left-hand structure, i.e., \mathbb{A} must be taken from some suitably defined class \mathbf{A} of structures, while \mathbb{B} is any arbitrary structure from the class “ $-$ ” of all finite structures.¹ Thus, we face the so-called *uniform* constraint satisfaction problem, shortly denoted as $\text{CSP}(\mathbf{A}, -)$, where both structures are part of the input (nothing is fixed).

The decision problem $\text{CSP}(\mathbf{A}, -)$ has intensively been studied in the literature, and various classes of structures over which it can be solved in polynomial time have already been singled out (see [6, 11, 18, 1], and the references therein). These approaches, called *decomposition methods*, are based on properties of the hypergraph $\mathcal{H}_{\mathbb{A}}$ associated with each structure $\mathbb{A} \in \mathbf{A}$. In fact, it is well-known that, for the class \mathbf{A}_a of all structures whose associated hypergraphs are acyclic, $\text{CSP}(\mathbf{A}_a, -)$ is efficiently solvable by just enforcing *generalized arc consistency* (GAC)—roughly, by filtering constraint relations until every pair of constraints having some variables \bar{X} in common agree on \bar{X} (that is, they have precisely the same set of allowed tuples of values on these variables \bar{X}).

Larger “islands of tractability” are then identified by generalizing hypergraph acyclicity. To this end, every decomposition method DM associates with any hypergraph $\mathcal{H}_{\mathbb{A}}$ some measure w of its cyclicity, called the *DM-width* of $\mathcal{H}_{\mathbb{A}}$. The tractable classes \mathbf{A} of instances (according to DM) are those (with

¹ Note that the finite property is a feature of this framework, and not a simplifying assumption. E.g., on structures with possibly infinite domains, the open question in [10] (just recently answered by [15] on finite structures) would have been solved in 1993 [23].

hypergraphs) having bounded width, that is, whose degree of cyclicity is below some fixed threshold. For every instance \mathbb{A} in such a class \mathbf{A} and every structure \mathbb{B} , the instance (\mathbb{A}, \mathbb{B}) can be solved in polynomial-time by exploiting the solutions of a set of suitable subproblems, that we call *views*, each one solvable in polynomial-time (in fact, exponential in the—fixed—width, for all known methods). In particular, the idea is to arrange some of these views in a tree, called *decomposition*, in order to exploit the known algorithms for acyclic instances. In fact, whenever such a tree exists, instances can be solved by just enforcing GAC on the available views, even without computing explicitly any decomposition. This very general approach traces back to the seminal database paper [10], and it is based on the graph-theoretic notion of *tree-projection* of the pair of hypergraphs $(\mathcal{H}_{\mathbb{A}}, \mathcal{H}_{\mathcal{V}})$, associated with the input structure \mathbb{A} and with the structure \mathcal{V} of the available views, respectively (tree projections are formally defined in Section 2).

For instance, assume that the fixed threshold on the width is k : in the *generalized hypertree-width* method [13], the available views are all subproblems involving at most k constraints from the given CSP instance; in the case of *treewidth* [21], the views are all subproblems involving at most k variables; for *fractional hypertree-width*, the views are all subproblems having fractional cover-width at most k (in fact, if we require that they are computable in polynomial-time, we may instead use those subproblems defined in [19] to compute a $O(k^3)$ approximation of this notion).

Note that, for the special case of generalized hypertree-width, the fact that enforcing GAC on all clusters of k constraints is sufficient to solve the given instance, without computing a decomposition, has been re-derived in [5] (with proof techniques different from those in [10]). Moreover, [5] actually provided a stronger result, as it is proved that this property holds even if there is some homomorphically equivalent subproblem having generalized hypertree-width at most k . However, the corresponding *only if* result is missing in that paper, and characterizing the precise power of this GAC procedure for the views obtained from all clusters of k constraints (short: k -GAC) remained an open question. For any class \mathbf{A} of instances having bounded arity (i.e., with a fixed maximum number of variables in any constraint scope of every instance of the class), the question has been answered in [2]: $\forall \mathbb{A} \in \mathbf{A}$, k -GAC is correct for every right-hand structure \mathbb{B} if, and only if, the core of \mathbb{A} has tree width at most k (recall that treewidth and generalized hypertree-width identify the same set of bounded-arity tractable classes). In its full version, the answer to this open question follows from a recent result in [15] (see Theorem 2-bis).

In fact, for any recursively enumerable class of bounded-arity structures \mathbf{A} , it is known that this method is essentially optimal: $\text{CSP}(\mathbf{A}, -)$ is solvable in polynomial time *if, and only if*, the cores of the structures in \mathbf{A} have bounded treewidth (under standard complexity theoretic assumptions) [17]. Note that the latter condition may be equivalently stated as follows: for every $\mathbb{A} \in \mathbf{A}$ there is some \mathbb{A}' homomorphically equivalent to \mathbb{A} and such that its treewidth is below the required fixed threshold. For short, we say that such a class has bounded treewidth modulo homomorphic equivalence.

Things with unbounded-arity classes are not that clear. Generalized hypertree-width does not characterize all classes of (arbitrary) structures where $\text{CSP}(\mathbf{A}, -)$ is solvable in polynomial time [18]. It seems that a useful characterization may be obtained by relaxing the typical requirement that views are computable in polynomial time, and by requiring instead that such tasks are fixed-parameter tractable (FPT) [9]. In fact, towards establishing such characterization, it was recently shown in [20] that (under some reasonable technical assumptions) the problem $\text{CSP}(\mathcal{H})$, i.e., $\text{CSP}(\mathbf{A}, -)$ restricted to the instances whose associated hypergraphs belong to the class \mathcal{H} , is FPT *if, and only if*, hypergraphs in \mathcal{H} have bounded *submodular* width—a new hypergraph measure more general than *fractional* hypertree-width and, hence, than generalized hypertree-width.

It is worthwhile noting that the above mentioned tractability results for classes of instances defined modulo homomorphically equivalence are actually tractability results for the *promise* version of the problem. In fact, unless $P = NP$, there is no polynomial-time algorithm that may check whether a given instance \mathbb{A} actually belongs to such a class \mathbf{A} . In particular, it has been observed by different authors [24, 4] that there are classes of instances having bounded treewidth modulo homomorphically equivalence

for which answers computable in polynomial time cannot be trusted. That is, unless $P = NP$, there is no efficient way to distinguish whether a “yes” answer means that there is some solution of the problem, or that $\mathbb{A} \notin \mathbf{A}$.

In this paper, besides promise problems, we also consider the so-called *no-promise* problems, which seem more appealing for practical applications. In this case, either certified solutions are computed, or the promise $\mathbb{A} \in \mathbf{A}$ is correctly disproved. For instance, the algorithm in [5] solves the no-promise search-problem of computing a homomorphism for a given CSP instance (\mathbb{A}, \mathbb{B}) . This algorithm either computes such a homomorphism or concludes that $\mathcal{H}_{\mathbb{A}}$ has generalized hypertree-width greater than k .

1.2 Enumeration Problems

While the structural tractability of deciding whether CSP instances admit solutions has been deeply studied in the literature, the structural tractability of the corresponding computation problem received considerably less attention so far [4], though this is certainly a more appealing problem for practical applications. In particular, it is well-known that for classes of CSPs where the decision problem is tractable and a self-reduction argument applies the enumeration problem is tractable too [8, 7]. Roughly, these classes have a suitable closure property such that one may fix values for the variables without going out of the class, and thus may solve the computation problem by using the (polynomial-time) algorithm for the decision problem as an oracle. In fact, for the non-uniform CSP problem, the tractability of the decision problem always entails the tractability of the search problem [7]. As observed above, this is rather different from what happens in the uniform CSP problem that we study in this paper, where this property does not hold (see [24, 4], and Proposition 1), and thus a specific study for the computation problem is meaningful and necessary.

In this paper, we embark on this study, by focusing on the problem ECSP of enumerating (possibly projected) solutions. Since even easy instances may have an exponential number of solutions, tractability means here having algorithms that compute solutions *with polynomial delay* (WPD): An algorithm M solves WPD a computation problem P if there is a polynomial $p(\cdot)$ such that, for every instance of P of size n , M discovers if there are no solutions in time $O(p(n))$; otherwise, it outputs all solutions in such a way that a new solution is computed within $O(p(n))$ time from the previous one.

Before stating our contribution, it is worthwhile noting that there are different facets of the enumeration problem, and thus different research directions to be explored:

(*Which Decomposition Methods?*) We considered the more general framework of the tree projections, where subproblems (views) may be completely arbitrary, so that our results are smoothly inherited by all (known) decomposition methods. We remark that this choice posed interesting technical challenges to our analysis, and called for solution approaches that were not explored in the earlier literature on traditional methods, such as treewidth. For instance, in this general context, we cannot speak anymore of “the core” of a structure, because different isomorphic cores may have different structural properties with respect to the available views.

(*Only full solutions or possibly projected solutions?*) In this paper, an ECSP instance is a triple $(\mathbb{A}, \mathbb{B}, O)$, for which we have to compute all solutions (homomorphisms) projected to a set of desired output variables O , denoted by $\mathbb{A}^{\mathbb{B}}[O]$. We believe this is the more natural approach. Indeed, modeling real-world applications through CSP instances typically requires the use of “auxiliary” variables, whose precise values in the solutions are not relevant for the user, and that are (usually) filtered-out from the output. In these cases, computing all combinations of their values occurring in solutions means wasting time, possibly exponential time. Of course, this aspect is irrelevant for the problem of computing just one solution, but is crucial for the enumeration problem.

(*Should classes of structures be aware of output variables?*) This is an important technical question. We are interested in identifying classes of tractable instances based on properties of their left-hand structures, while right-hand structures have no restrictions. What about output variables? In principle,

structural properties may or may not consider the possible output variables, and in fact both approaches have been explored in the literature (see, e.g., [17]), and both approaches are dealt with in this paper. In the former output-aware case, possible output variables are suitably described in the instance structure. Unlike previous approaches that considered additional “virtual” constraints covering together all possible output variables [17], in this paper possible output variables are described as those variables X having a domain constraint $dom(X)$, that is, a distinguished unary constraint specifying the domain of this variable. Such variables are said domain restricted. In fact, this choice reflects the classical approach in constraint satisfaction systems, where variables are typically associated with domains, which are heavily exploited by constraint propagation algorithms. Note that this approach does not limit the number of solutions, while in the tractable classes considered in [17] only instances with a polynomial number of (projected) solutions may be dealt with. As far as the latter case of arbitrary sets of output variables is considered, observe that in general stronger conditions are expected to be needed for tractability. Intuitively, since we may focus on any desired substructure, no strange situations may occur, and the full instance should be really tractable.

1.3 Contribution

Output-aware classes of ECSPs:

- (1) We define a property for pairs (\mathbb{A}, O) , where \mathbb{A} is a structure and $O \subseteq A$ is a set of variables, that allows us to characterize the classes of tractable instances. Roughly, we say that (\mathbb{A}, O) is *tp-covered* through the decomposition method DM if variables in O occur in a tree projection of a certain hypergraph w.r.t. to the (hypergraph associated with the) views defined according to DM.
- (2) We describe an algorithm that solves the promise enumeration problem, by computing with polynomial delay all solutions of a given instance $(\mathbb{A}, \mathbb{B}, O)$, whenever (\mathbb{A}, O) is *tp-covered* through DM.
- (3) For the special case of (generalized hyper)tree width, we show that the above condition is also necessary for the correctness of the proposed algorithm (for every \mathbb{B}). In fact, for these traditional decomposition methods we now have a complete characterization of the power of the k -GAC approach.
- (4) For recursively enumerable classes of structures having bounded arity, we exhibit a dichotomy showing that the above tractability result is tight, for DM = treewidth (and assuming $FPT \neq W[1]$).

ECSP instances over arbitrary output variables:

- (1) We describe an algorithm that, on input $(\mathbb{A}, \mathbb{B}, O)$, solves the no-promise enumeration problem. In particular, either all solutions are computed, or it infers that there exists no tree projection of $\mathcal{H}_{\mathbb{A}}$ w.r.t. $\mathcal{H}_{\mathbb{B}}$ (the hypergraph associated with the views defined according to DM). This algorithm generalizes to the tree projection framework the enumeration algorithm of projected solutions recently proposed for the special case of treewidth [4].
- (2) Finally, we give some evidence that, for bounded arity classes of instances, we cannot do better than this. In particular, having bounded width tree-decompositions of the full structure seems a necessary condition for enumerating WPD. We speak of “evidence,” instead of saying that our result completely answers the open question in [17, 4], because our dichotomy theorem focuses on classes of structures satisfying the technical property of being closed under taking minors (in fact, the same property assumed in the first dichotomy result on the complexity of the decision problem on classes of graphs [16]).

2 Preliminaries: Relational Structures and Homomorphisms

A constraint satisfaction problem may be formalized as a relational homomorphism problem. A vocabulary τ is a finite set of relation symbols of specified arities. A relational structure \mathbb{A} over τ consists of a universe A and an r -ary relation $R^{\mathbb{A}} \subseteq A^r$, for each relation symbol R in τ .

If \mathbb{A} and \mathbb{A}' are two relational structures over disjoint vocabularies, we denote by $\mathbb{A} \uplus \mathbb{A}'$ the relational structure over the (disjoint) union of their vocabularies, whose domain (resp., set of relations) is the union of those of \mathbb{A} and \mathbb{A}' .

A *homomorphism* from a relational structure \mathbb{A} to a relational structure \mathbb{B} is a mapping $h : A \mapsto B$ such that, for every relation symbol $R \in \mathbb{A}$, and for every tuple $\langle a_1, \dots, a_r \rangle \in R^{\mathbb{A}}$, it holds that $\langle h(a_1), \dots, h(a_r) \rangle \in R^{\mathbb{B}}$. For any set $X \subseteq A$, denote by $h[X]$ the restriction of h to X . The set of all possible homomorphisms from \mathbb{A} to \mathbb{B} is denoted by $\mathbb{A}^{\mathbb{B}}$, while $\mathbb{A}^{\mathbb{B}}[X]$ denotes the set of their restrictions to X .

An instance of the constraint satisfaction problem (CSP) is a pair (\mathbb{A}, \mathbb{B}) where \mathbb{A} is called a *left-hand structure* (short: ℓ -structure) and \mathbb{B} is called a *right-hand structure* (short: r -structure). In the classical decision problem, we have to decide whether there is a homomorphism from \mathbb{A} to \mathbb{B} , i.e., whether $\mathbb{A}^{\mathbb{B}} \neq \emptyset$. In an instance of the corresponding enumeration problem (denoted by ECSP) we are additionally given a set of *output* elements $O \subseteq A$; thus, an instance has the form $(\mathbb{A}, \mathbb{B}, O)$. The goal is to compute the restrictions to O of all possible homomorphisms from \mathbb{A} to \mathbb{B} , i.e., $\mathbb{A}^{\mathbb{B}}[O]$. If $O = \emptyset$, the computation problem degenerates to the decision one. Formally, let $h_\phi : \emptyset \mapsto \text{true}$ denote (the constant mapping to) the Boolean value *true*; then, define $\mathbb{A}^{\mathbb{B}}[\emptyset] = \{h_\phi\}$ (resp., $\mathbb{A}^{\mathbb{B}}[\emptyset] = \emptyset$) if there is some (resp., there is no) homomorphism from \mathbb{A} to \mathbb{B} .

In the constraint satisfaction jargon, the elements of A (the domain of the ℓ -structure \mathbb{A}) are the variables, and there is a constraint $C = (\langle a_1 \dots, a_r \rangle, R^{\mathbb{B}})$ for every tuple $\langle a_1 \dots, a_r \rangle \in R^{\mathbb{A}}$ and every relation symbol $R \in \tau$. The tuple of variables is usually called the *scope* of C , while $R^{\mathbb{B}}$ is called the relation of C . Any homomorphism from \mathbb{A} to \mathbb{B} is thus a mapping from the variables in A to the elements in B (often called domain values) that satisfies all constraints, and it is also called a *solution* (or a *projected solution*, if it is restricted to a subset of the variables).

Two relational structures \mathbb{A} and \mathbb{A}' are homomorphically equivalent if there is a homomorphism from \mathbb{A} to \mathbb{A}' and vice-versa. A structure \mathbb{A}' is a *substructure* of \mathbb{A} if $A' \subseteq A$ and $R^{\mathbb{A}'} \subseteq R^{\mathbb{A}}$, for each symbol $R \in \tau$. Moreover, \mathbb{A}' is a *core* of \mathbb{A} if it is a substructure of \mathbb{A} such that: (1) there is a homomorphism from \mathbb{A} to \mathbb{A}' , and (2) there is no substructure \mathbb{A}'' of \mathbb{A}' , with $\mathbb{A}'' \neq \mathbb{A}'$, satisfying (1).

3 Decomposition Methods, Views, and Tree Projections

Throughout the following sections we assume that (\mathbb{A}, \mathbb{B}) is a given connected CSP instance, and we shall seek to compute its solutions (possibly restricted to a desired set of output variables) by combining the solutions of suitable sets of subproblems, available as additional distinguished constraints called *views*.

Let \mathbb{A}_V be an ℓ -structure with the same domain as \mathbb{A} . We say that \mathbb{A}_V is a *view structure* (short: *v-structure*) if

- its vocabulary τ_V is disjoint from the vocabulary τ of \mathbb{A} ;
- every relation $R^{\mathbb{A}_V}$ contains a single tuple whose variables will be denoted by $\text{var}(R^{\mathbb{A}_V})$. That is, there is a one-to-one correspondence between views and relation symbols in τ_V , so that we shall often use the two terms interchangeably;
- for every relation $R \in \tau$ and every tuple $t \in R^{\mathbb{A}}$, there is some relation $R_t \in \tau_V$, called *base view*, such that $\{t\} = R_t^{\mathbb{A}_V}$, i.e., for every constraint in \mathbb{A} there is a corresponding view in \mathbb{A}_V .

Let \mathbb{B}_V be an r -structure. We say that \mathbb{B}_V is *legal* (w.r.t. \mathbb{A}_V and (\mathbb{A}, \mathbb{B})) if

- its vocabulary is τ_V ;




Fig. 1. A structure \mathbb{A} . A hypergraph $\mathcal{H}_{\mathbb{A}_V}$ such that $(\mathcal{H}_{\mathbb{A}}, \mathcal{H}_{\mathbb{A}_V})$ has no tree projections. Two hypergraphs $\mathcal{H}_{\mathbb{A}'}$ and $\mathcal{H}_{\mathbb{A}''}$, where \mathbb{A}' and \mathbb{A}'' are cores of \mathbb{A} . A tree projection \mathcal{H}_a of $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\mathbb{A}_V})$.

- For every view $R \in \tau_V$, $R^{\mathbb{B}_V} \supseteq \mathbb{A}^{\mathbb{B}}[w]$ holds, where $w = \text{var}(R^{\mathbb{A}_V})$. That is, every subproblem is not more restrictive than the full problem.
- For every base view $R_t \in \tau_V$, $R_t^{\mathbb{B}_V} \subseteq R^{\mathbb{B}}$. That is, any base view is at least as restrictive as the “original” constraint associated with it.

The following fact immediately follows from the above properties.

Fact 1 Let \mathbb{B}_V be any r -structure that is legal w.r.t. \mathbb{A}_V and (\mathbb{A}, \mathbb{B}) . Then, $\forall O \subseteq A$, the ECSP instance $(\mathbb{A}_V, \mathbb{B}_V, O)$ has the same set of solutions as $(\mathbb{A}, \mathbb{B}, O)$.

In fact, all structural decomposition methods define some way to build the views to be exploited for solving the given CSP instance. In our framework, we associate with any decomposition method DM a pair of polynomial-time computable functions ℓ -DM and r -DM that, given any CSP instance (\mathbb{A}, \mathbb{B}) , compute the pair $(\mathbb{A}_V, \mathbb{B}_V)$, where $\mathbb{A}_V = \ell\text{-DM}(\mathbb{A})$ is a v -structure, and $\mathbb{B}_V = r\text{-DM}(\mathbb{A}, \mathbb{B})$ is a legal r -structure.²

For instance, for any fixed natural number k , the *generalized hypertree decomposition method* [12] (short: hw_k) is associated with the functions $\ell\text{-}hw_k$ and $r\text{-}hw_k$ that, given a CSP instance (\mathbb{A}, \mathbb{B}) , build the pair $(\ell\text{-}hw_k(\mathbb{A}), r\text{-}hw_k(\mathbb{A}, \mathbb{B}))$ where, for each subset C of at most k constraints from (\mathbb{A}, \mathbb{B}) , there is a view R_C such that: (1) $\text{var}(R_C^{\ell\text{-}hw_k(\mathbb{A})})$ is the set of all variables occurring in C , and (2) the tuples in $R_C^{r\text{-}hw_k(\mathbb{A}, \mathbb{B})}$ are the solutions of the subproblem encoded by C . Similarly, the *tree decomposition method* [21] (tw_k) is defined as above, but we consider each subset of at most k variables in \mathbb{A} instead of each subset of at most k constraints.

3.1 Tree Projections for CSP Instances

In this paper we are interested in restrictions imposed on left-hand structures of CSP instances, based on some decomposition method DM. To this end, we associate with any ℓ -structure \mathbb{A} a hypergraph $\mathcal{H}_{\mathbb{A}} = (A, H)$, whose set of nodes is equal to the set of variables A and where, for each constraint scope in $R^{\mathbb{A}}$, the set H of hyperedges contains a hyperedge including all its variables (no further hyperedge is in H). In particular, the v -structure $\mathbb{A}_V = \ell\text{-DM}(\mathbb{A})$ is associated with a hypergraph $\mathcal{H}_{\mathbb{A}_V} = (A, H)$, whose set of nodes is the set of variables A and where, for each view $R \in \tau_V$, the set H contains

² A natural extension of this notion we may be to consider FPT decomposition methods, where functions ℓ -DM and r -DM are computable in fixed-parameter polynomial-time. For the sake of presentation and of space, we do not consider FPT decomposition methods in this paper, but our results can be extended to them rather easily.

the hyperedge $\text{var}(R^{\mathbb{A}_v})$. In the following, for any hypergraph \mathcal{H} , we denote its nodes and edges by $\text{nodes}(\mathcal{H})$ and $\text{edges}(\mathcal{H})$, respectively.

Example 1. Consider the ℓ -structure \mathbb{A} whose vocabulary just contains the binary relation symbol R , and such that $R^{\mathbb{A}} = \{\langle F, E \rangle, \langle A, E \rangle, \langle A, C \rangle, \langle A, B \rangle, \langle B, C \rangle, \langle D, B \rangle, \langle D, C \rangle\}$. Such a simple one-binary-relation structure may be easily represented by the directed graph in the left part of Figure 1, where edge orientation reflects the position of the variables in R . In this example, the associated hypergraph $\mathcal{H}_{\mathbb{A}}$ is just the undirected version of this graph. Let DM be a method that, on input \mathbb{A} , builds the v -structure $\mathbb{A}_v = \ell\text{-DM}(\mathbb{A})$ consisting of the seven base views of the form R_t , for each tuple $t \in R^{\mathbb{A}}$, plus the three relations R_1, R_2 , and R_3 such that $R_1^{\mathbb{A}_v} = \{\langle A, E, F \rangle\}$, $R_2^{\mathbb{A}_v} = \{\langle A, B, C, F \rangle\}$, and $R_3^{\mathbb{A}_v} = \{\langle C, D, F \rangle\}$. Figure 1 also reports $\mathcal{H}_{\mathbb{A}_v}$. \triangleleft

A hypergraph \mathcal{H} is *acyclic* iff it has a *join tree* [3], i.e., a tree $\mathcal{JT}(\mathcal{H})$, whose vertices are the hyperedges of \mathcal{H} , such that if a node X occurs in two hyperedges h_1 and h_2 of \mathcal{H} , then h_1 and h_2 are connected in $\mathcal{JT}(\mathcal{H})$, and X occurs in each vertex on the unique path linking h_1 and h_2 in $\mathcal{JT}(\mathcal{H})$.

For two hypergraphs \mathcal{H}_1 and \mathcal{H}_2 , we write $\mathcal{H}_1 \leq \mathcal{H}_2$ iff each hyperedge of \mathcal{H}_1 is contained in at least one hyperedge of \mathcal{H}_2 . Let $\mathcal{H}_1 \leq \mathcal{H}_2$. Then, a *tree projection* of \mathcal{H}_1 with respect to \mathcal{H}_2 is an acyclic hypergraph \mathcal{H}_a such that $\mathcal{H}_1 \leq \mathcal{H}_a \leq \mathcal{H}_2$. Whenever such a hypergraph \mathcal{H}_a exists, we say that the pair $(\mathcal{H}_1, \mathcal{H}_2)$ has a tree projection (also, we say that \mathcal{H}_1 has a tree projection w.r.t. \mathcal{H}_2). The problem of deciding whether a pair of hypergraphs has a tree projection is called the *tree projection problem*, and it has recently been proven to be NP-complete [14].

Example 2. Consider again the setting of Example 1. It is immediate to check that the pair of hypergraphs $(\mathcal{H}_{\mathbb{A}}, \mathcal{H}_{\mathbb{A}_v})$ does not have any tree projection. Consider instead the (hyper)graph $\mathcal{H}_{\mathbb{A}'}$ reported on the right of Figure 1. The acyclic hypergraph \mathcal{H}_a is a tree projection of $\mathcal{H}_{\mathbb{A}'}$ w.r.t. $\mathcal{H}_{\mathbb{A}_v}$. In particular, note that the hyperedge $\{A, B, C\} \in \text{edges}(\mathcal{H}_a)$ “absorbs” the cycle in $\mathcal{H}_{\mathbb{A}'}$, and that $\{A, B, C\}$ is in its turn contained in the hyperedge $\{A, B, C, F\} \in \text{edges}(\mathcal{H}_{\mathbb{A}_v})$. \triangleleft

Note that all the (known) structural decomposition methods can be recast as special cases of tree projections, since they just differ in how they define the set of views to be built for evaluating the CSP instance. For instance, a hypergraph $\mathcal{H}_{\mathbb{A}}$ has generalized hypertree width (resp., treewidth) at most k if and only if there is a tree projection of $\mathcal{H}_{\mathbb{A}}$ w.r.t. $\mathcal{H}_{\ell\text{-}hw_k(\mathbb{A})}$ (resp., w.r.t. $\mathcal{H}_{\ell\text{-}tw_k(\mathbb{A})}$).

However, the setting of tree projections is more general than such traditional decomposition approaches, as it allows us to consider arbitrary sets of views, which often require more care and different techniques. As an example, we shall illustrate below that in the setting of tree projections it does not make sense to talk about “the” core of an ℓ -structure, because different isomorphic cores may differently behave with respect to the available views. This phenomenon does not occur, e.g., for generalized hypertree decompositions, where all combinations of k constraints are available as views.

Example 3. Consider the structure \mathbb{A} illustrated in Example 1, and the structures \mathbb{A}' and \mathbb{A}'' over the same vocabulary as \mathbb{A} , and such that $R^{\mathbb{A}'} = \{\langle A, C \rangle, \langle A, B \rangle, \langle B, C \rangle\}$ and $R^{\mathbb{A}''} = \{\langle B, C \rangle, \langle D, B \rangle, \langle D, C \rangle\}$. The hypergraphs $\mathcal{H}_{\mathbb{A}'}$ and $\mathcal{H}_{\mathbb{A}''}$ are reported in Figure 1. Note that \mathbb{A}' and \mathbb{A}'' are two (isomorphic) cores of \mathbb{A} , but they have completely different structural properties. Indeed, $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\mathbb{A}_v})$ admits a tree projection (recall Example 2), while $(\mathcal{H}_{\mathbb{A}''}, \mathcal{H}_{\mathbb{A}_v})$ does not. \triangleleft

3.2 CSP Instances and tp-Coverings

We complete the picture of our unifying framework to deal with decomposition methods for constraint satisfaction problems, by illustrating some recent results in [15], which will be useful to our ends. Let us start by stating some preliminary definitions.

For a set of variables $O = \{X_1, \dots, X_r\}$, let \mathbb{S}_O denote the structure with a fresh r -ary relation symbol R_O and domain O , such that $R_O^{\mathbb{S}_O} = \{\langle X_1, \dots, X_r \rangle\}$.

Definition 1. Let \mathbb{A}_V be a v -structure. A set of variables $O \subseteq A$ is *tp-covered* in \mathbb{A}_V if there exists a core \mathbb{A}' of $\mathbb{A} \uplus \mathbb{S}_O$ such that $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\mathbb{A}_V})$ has a tree projection.³ \square

For instance, it is easily seen that the variables $\{A, B, C\}$ are *tp-covered* in the v -structure \mathbb{A}_V discussed in Example 1. In particular, note that the structure $\mathbb{A} \uplus \mathbb{S}_{\{A, B, C\}}$ is associated with the same hypergraph $\mathcal{H}_{\mathbb{A}'}$ that has a tree projection w.r.t. $\mathcal{H}_{\mathbb{A}_V}$ (cf. Example 3). Instead, the variables $\{B, C, D\}$ are not *tp-covered* in \mathbb{A}_V .

Given a CSP instance $(\mathbb{A}_V, \mathbb{B}_V)$, we denote by $\text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$ the r -structure that is obtained by enforcing generalized arc consistency on $(\mathbb{A}_V, \mathbb{B}_V)$.

The following result, proved in [15] for a different setting, states the precise relationship between generalized-arc-consistent views and *tp-covered* sets of variables.

Theorem 2. Let \mathbb{A} be an ℓ -structure, and let \mathbb{A}_V be a v -structure. The following are equivalent:

- (1) A set of variables $O \subseteq A$ is *tp-covered* in \mathbb{A}_V ;
- (2) For every r -structure \mathbb{B} , for every r -structure \mathbb{B}_V that is legal w.r.t. \mathbb{A}_V and (\mathbb{A}, \mathbb{B}) , and for every relation $R \in \tau_V$ with $O \subseteq \text{var}(R^{\mathbb{A}_V})$, $R^{\text{GAC}(\mathbb{A}_V, \mathbb{B}_V)}[O] = \mathbb{A}^{\mathbb{B}}[O]$.

Note that the result answered a long standing open question [10, 23] about the relationship between the existence of tree projections and (local and global) consistency properties in databases [15]. In words, the result states that just enforcing generalized arc consistency on the available views is a sound and complete procedure to solve ECSP instances *if, and only if*, we are interested in (projected) solutions over output variables that are *tp-covered* and occur together in some available view. Thus, in these cases, all solutions can be computed in polynomial time. The more general case where output variables are arbitrary (i.e., not necessarily included in some available view) is explored in the rest of this paper.

We now leave the section by noticing that as a consequence of Theorem 2, we can characterize the power of local-consistency for any decomposition method DM such that, for each pair (\mathbb{A}, \mathbb{B}) , each view in $\mathbb{B}_V = r\text{-DM}(\mathbb{A}, \mathbb{B})$ contains the solutions of the subproblem encoded by the constraints over which it is defined. For the sake of simplicity, we state below the result specialized to the well-known decomposition methods tw_k and hw_k .

Theorem 2-bis. Let DM be a decomposition method in $\{tw_k, hw_k\}$, let \mathbb{A} be an ℓ -structure, and let $\mathbb{A}_V = \ell\text{-DM}(\mathbb{A})$. The following are equivalent:

- (1) A set of variables $O \subseteq A$ is *tp-covered* in \mathbb{A}_V ;
- (2) For every r -structure \mathbb{B} , and for every relation $R \in \tau_V$ with $O \subseteq \text{var}(R^{\mathbb{A}_V})$, $R^{\text{GAC}(\mathbb{A}_V, \mathbb{B}_V)}[O] = \mathbb{A}^{\mathbb{B}}[O]$, where $\mathbb{B}_V = r\text{-DM}(\mathbb{A}, \mathbb{B})$.

Proof (Sketch). Preliminarily, it is easy to see that (2) in Theorem 2 may be equivalently stated as follows:

- (2') For every r -structure \mathbb{B} , for every r -structure \mathbb{B}_V that is legal w.r.t. \mathbb{A}_V and (\mathbb{A}, \mathbb{B}) and such that $\mathbb{B}_V = \text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$, and for every relation $R \in \tau_V$ with $O \subseteq \text{var}(R^{\mathbb{A}_V})$, $R^{\mathbb{B}_V}[O] = \mathbb{A}^{\mathbb{B}}[O]$.

The fact that (1) \Rightarrow (2) trivially follows from Theorem 2. We have to show that (2) \Rightarrow (1) holds as well. To this end, observe that if O is not *tp-covered* in $\ell\text{-DM}(\mathbb{A})$, by Theorem 2 (actually, (1) \Rightarrow (2')), we can conclude the existence of: (1) an r -structure \mathbb{B} , (2) an r -structure \mathbb{B}_V that is legal w.r.t. \mathbb{A}_V and (\mathbb{A}, \mathbb{B}) and such that $\mathbb{B}_V = \text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$, and (3) a relation $R \in \tau_V$ with $O \subseteq \text{var}(R^{\mathbb{A}_V})$ such that $R^{\mathbb{B}_V}[O] \neq \mathbb{A}^{\mathbb{B}}[O]$ (of course, $R^{\mathbb{B}_V}[O] \supset \mathbb{A}^{\mathbb{B}}[O]$ by the legality of \mathbb{B}_V). Consider now the r -structure

³ For the sake of completeness, note that we use here a core \mathbb{A}' because we found it more convenient for the presentation and the proofs. However, it is straightforward to check that this notion can be equivalently stated in terms of any structure homomorphically equivalent to $\mathbb{A} \uplus \mathbb{S}_O$. The same holds for the related Definition 3.

$\mathbb{B}'_{\mathcal{V}} = r\text{-}hw_k(\mathbb{A}, \mathbb{B})$. Recall that each view in $\mathbb{B}'_{\mathcal{V}}$ contains *all* the solutions of the subproblem encoded by the constraints over which it is defined. Since $\mathbb{B}_{\mathcal{V}} = \text{GAC}(\mathbb{A}_{\mathcal{V}}, \mathbb{B}_{\mathcal{V}})$, it can be shown that each view in $\mathbb{B}_{\mathcal{V}}$ contains *only* solutions of the subproblem encoded by the constraints over which it is defined. Thus, for each relation $R \in \tau_{\mathcal{V}}$, $R^{\mathbb{B}'_{\mathcal{V}}} \supseteq R^{\mathbb{B}_{\mathcal{V}}}$ holds, which implies $R^{\mathbb{B}'_{\mathcal{V}}}[O] \supset \mathbb{A}^{\mathbb{B}}[O]$. The same line of reasoning applies to the tree decomposition method. \triangleleft

Note that if we consider decision problem instances ($O = \emptyset$) and the treewidth method ($\mathbb{A}_{\mathcal{V}} = \ell\text{-}tw_k(\mathbb{A})$), from Theorem 2-bis, we (re-)obtain the nice characterization of [2] about the relationship between k -local consistency and treewidth modulo homomorphic equivalence. If we consider generalized hypertree-width ($\mathbb{A}_{\mathcal{V}} = \ell\text{-}hw_k(\mathbb{A})$), we get the answer to the corresponding open question for the unbounded arity case, that is, the precise power of the procedure enforcing k -union (of constraints) consistency (i.e., the power of the algorithm for the decision problem described in [5]).

4 Enumerating Solutions of Output-Aware CSP Instances

The goal of this section is to study the problem of enumerating CSP solutions for classes of instances where possible output variables are part of the structure of the given instance. This is formalized by assuming that the relational structure contains *domain constraints* that specify the domains for such variables.

Definition 2. A variable $X \in A$ is *domain restricted* in the ℓ -structure \mathbb{A} if there exists a unary distinguished (domain) relation symbol $\text{dom}(X) \in \tau$ such that $\{\langle X \rangle\} = \text{dom}(X)^{\mathbb{A}}$. The set of all domain restricted variables is denoted by $\text{drv}(\mathbb{A})$. \square

We say that an ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ is domain restricted if $O \subseteq \text{drv}(\mathbb{A})$. Of course, if it is not, then one may easily build in linear time an equivalent domain-restricted ECSP instance where an additional fresh unary constraint is added for every output variable, whose values are taken from any constraint relation where that variable occurs. We say that such an instance is a domain-restricted version of $(\mathbb{A}, \mathbb{B}, O)$.

Figure 2 shows an algorithm, named $\text{ComputeAllSolutions}_{\text{DM}}$, that computes the solutions of a given ECSP instance. The algorithm is parametric w.r.t. any chosen decomposition method DM, and works as follows. Firstly, $\text{ComputeAllSolutions}_{\text{DM}}$ starts by transforming the instance $(\mathbb{A}, \mathbb{B}, O)$ into a domain restricted one, and by constructing the views in $(\mathbb{A}_{\mathcal{V}}, \mathbb{B}_{\mathcal{V}})$ via DM. Then, it invokes the procedure *Propagate*. This procedure backtracks over the output variables $\{X_1, \dots, X_m\}$: At each step i , it tries to assign a value to X_i from its domain view,⁴ and defines this value as the unique one available in that domain, in order to “propagate” such an assignment over all other views. This is accomplished by enforcing generalized arc-consistency each time the procedure is invoked. Eventually, whenever an assignment is computed for all the variables in O , this solution is returned in output, and the algorithm proceeds by backtracking again trying different values.

4.1 Tight Characterizations for the Correctness of $\text{ComputeAllSolutions}_{\text{DM}}$

To characterize the correctness of $\text{ComputeAllSolutions}_{\text{DM}}$, we need to define a structural property that is related to the one stated in Definition 1. Below, differently from Definition 1 where the set of output variables O is treated as a whole, each variable in O has to be tp-covered as a singleton set.

Definition 3. Let $(\mathbb{A}, \mathbb{B}, O)$ be an ECSP instance. We say that (\mathbb{A}, O) is *tp-covered* through DM if there is a core \mathbb{A}' of $\mathbb{A} \uplus \biguplus_{X \in O} \mathbb{S}_{\{X\}}$ such that $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\ell\text{-DM}(\mathbb{A})})$ has a tree projection. \square

⁴ With an abuse of notation, in the algorithm we denote by $\text{dom}(X)$ the base view in $\tau_{\mathcal{V}}$ associated with the input constraint $\text{dom}(X) \in \tau$ (in fact, no confusion may arise because the algorithm only works on views).

<p>Input: An ECSP instance $(\mathbb{A}, \mathbb{B}, O)$, where $O = \{X_1, \dots, X_m\}$;</p> <p>Output: $\mathbb{A}^{\mathbb{B}}[O]$;</p> <p>Method: update $(\mathbb{A}, \mathbb{B}, O)$ with any of its domain-restricted versions; let $\mathbb{A}_V := \ell\text{-DM}(\mathbb{A})$, $\mathbb{B}_V := r\text{-DM}(\mathbb{A}, \mathbb{B})$; invoke Propagate(1, $(\mathbb{A}_V, \mathbb{B}_V)$, m, $\langle \rangle$);</p> <hr/> <p>Procedure Propagate(i: integer, $(\mathbb{A}_V, \mathbb{B}_V)$: pair of structures, m: integer, $\langle a_1, \dots, a_{i-1} \rangle$: tuple of values in A^i);</p> <p>begin</p> <ol style="list-style-type: none"> 1. let $\mathbb{B}'_V := \text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$; 2. let $\text{activeValues} := \text{dom}(X_i)^{\mathbb{B}'_V}$; 3. for each element $\langle a_i \rangle \in \text{activeValues}$ do 4. if $i = m$ then 5. output $\langle a_1, \dots, a_{m-1}, a_m \rangle$; 6. else 7. update $\text{dom}(X_i)^{\mathbb{B}'_V}$ with $\{\langle a_i \rangle\}$; /* X_i is fixed to value a_i */ 8. Propagate($i + 1$, $(\mathbb{A}_V, \mathbb{B}'_V)$, m, $\langle a_1, \dots, a_{i-1}, a_i \rangle$); <p>end.</p>
--

Fig. 2. Algorithm $\text{ComputeAllSolutions}_{\text{DM}}$.

Note that the above definition is purely structural, because (the right-hand structure) \mathbb{B} plays no role there. In fact, we next show that this definition captures classes of instances where $\text{ComputeAllSolutions}_{\text{DM}}$ is correct.

Theorem 3. Let DM be a decomposition method, let \mathbb{A} be an ℓ -structure, and let $O \subseteq A$ be a set of variables. Assume that (\mathbb{A}, O) is tp -covered through DM . Then, for every r -structure \mathbb{B} , $\text{ComputeAllSolutions}_{\text{DM}}$ computes the set $\mathbb{A}^{\mathbb{B}}[O]$.

Proof (Sketch). Let \mathbb{B}_{in} be any r -structure. Preliminarily observe that if the original input instance, say $I_{in} = (\mathbb{A}_{in}, \mathbb{B}_{in}, O)$, is tp -covered through DM , the same property is enjoyed by its equivalent domain-restricted version, say $I_0 = (\mathbb{A}, \mathbb{B}_0, O)$, computed in the starting phase of the algorithm. Thus, there is a core \mathbb{A}' of $\mathbb{A} \uplus \biguplus_{X \in O} \mathbb{S}_{\{X\}}$ such that $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\mathbb{A}_V})$ has a tree projection, where $\mathbb{A}_V = \ell\text{-DM}(\mathbb{A})$. This entails that, $\forall X \in O$, $\{X\}$ is tp -covered in $\mathcal{H}_{\mathbb{A}_V}$. It is sufficient to show that, if $\mathbb{A}_0^{\mathbb{B}} \neq \emptyset$, at the generic call of **Propagate** with i as its first argument, activeValues is initialized at Step 2 with a non-empty set that contains all those values that $X_i \in O$ may take, in any solution of $(\mathbb{A}, \mathbb{B}_0)$ extending the current partial solution $\langle a_1, \dots, a_{i-1} \rangle$; otherwise ($\mathbb{A}_0^{\mathbb{B}} = \emptyset$), $\text{activeValues} = \emptyset$, and the algorithm correctly terminates with an empty output without ever entering the **for** cycle. For the sake of presentation, we just prove what happens in the first call. The generalization to the generic case is then straightforward.

Let $i = 1$ and assume that $\mathbb{B}'_V := \text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$ has been computed. From the tp -covered property of variables in O and Theorem 2, it follows that $\forall X \in O$, its domain view $\text{dom}(X)$ is such that $\text{dom}(X)^{\mathbb{B}'_V} = \mathbb{A}_0^{\mathbb{B}}[\{X\}]$. Thus, all values in the domain views associated with output variables occur in some solutions. This holds in particular for $\text{dom}(X_1)$ that is empty if, and only if, $\mathbb{A}_0^{\mathbb{B}} = \emptyset$, in which case the cycle is skipped and the algorithm immediately halts with an empty output. Assume now that this is not the case, so that $\text{activeValues} = \mathbb{A}_0^{\mathbb{B}}[\{X_1\}] \neq \emptyset$, and let a_1 be the chosen value at Step 3. Consider a new instance $I_1 = (\mathbb{A}, \mathbb{B}_1, O)$ where the domain constraint for X_1 contains the one value a_1 . From the above discussion it follows that $\mathbb{A}_1^{\mathbb{B}} \neq \emptyset$, and clearly the solutions of I_1 are all and only those of I_0 that extend the partial solution $\langle a_1 \rangle$. Moreover, it is easy to check that the r -structure \mathbb{B}'_V obtained after the execution of Step 7 is legal w.r.t. $(\mathbb{A}, \mathbb{B}_1)$, and recall that nothing is changed in the pair (\mathbb{A}, O) , which is (still) tp -covered through DM . Therefore, when we call recursively call **Propagate** at Step 8 with $i = 2$, we are in the same situation as in the first call, but going to enumerate the solutions of I_1 .

At the end of this call, we just repeat this procedure with the next available value for X_1 , say a_2 , until all elements in $activeValues = \mathbb{A}_0^{\mathbb{B}}[\{X_1\}]$ have been considered (and propagated). \triangleleft

We now complete the picture by observing that Definition 3 also provides the necessary conditions for the correctness of $ComputeAllSolutions_{DM}$. As in Theorem 2-bis, we state below the result specialized to the methods tw_k and hw_k .

Theorem 4. *Let DM be a decomposition method in $\{tw_k, hw_k\}$, let \mathbb{A} be an ℓ -structure, and let $O \subseteq A$ be a set of variables. Assume that, for every r -structure \mathbb{B} , $ComputeAllSolutions_{DM}$ computes $\mathbb{A}^{\mathbb{B}}[O]$. Then, (\mathbb{A}, O) is tp -covered through DM .*

Proof (Sketch). Assume that (\mathbb{A}, O) is not tp -covered through DM , and let $O' \subseteq O$ be a maximal set of output variables such that (\mathbb{A}, O') is tp -covered through DM . In the case where $O' = \emptyset$, there is no core \mathbb{A}' of \mathbb{A} such that $(\mathcal{H}_{\mathbb{A}'}, \mathcal{H}_{\ell-DM(\mathbb{A})})$ has a tree projection. Thus, we can apply Theorem 2-bis and conclude that there are an r -structure \mathbb{B} , and a relation $R \in \tau_V$ such that $\mathbb{A}^{\mathbb{B}}$ has a solution while $R^{GAC(\ell-DM(\mathbb{A}), \mathbb{B}_V)}$ is empty, with $\mathbb{B}_V = r-DM(\mathbb{A}, \mathbb{B})$. It follows that $ComputeAllSolutions_{DM}$ will not produce any output. Consider now the case where $O' \neq \emptyset$, and where any $X \in O \setminus O'$ is a variable such that $(\mathbb{A}, O' \cup \{X\})$ is not tp -covered through DM . Let $\tilde{\mathbb{A}}$ be the relational structure $\mathbb{A} \uplus \biguplus_{Y \in O'} \mathbb{S}_{\{Y\}}$, which is such that $(\mathcal{H}_{\tilde{\mathbb{A}}}, \mathcal{H}_{\ell-DM(\mathbb{A})})$ has a tree projection. Then, $\{X\}$ is not tp -covered in $\ell-DM(\tilde{\mathbb{A}})$. By Theorem 2-bis, there are an r -structure \mathbb{B} , and a relation $R \in \tau_V$ with $\{X\} \subseteq var(R^{\tilde{\mathbb{A}}_V})$ such that $R^{GAC(\tilde{\mathbb{A}}_V, \mathbb{B}_V)}[\{X\}] \supset \tilde{\mathbb{A}}^{\mathbb{B}}[\{X\}]$, where $\mathbb{B}_V = r-DM(\tilde{\mathbb{A}}, \mathbb{B})$. In fact, we can show that such a “counterexample” structure \mathbb{B} can be chosen in such a way that there are (full) solutions h for the problem having the following property: some values in $R^{GAC(\tilde{\mathbb{A}}_V, \mathbb{B}_V)}[\{X\}] \setminus \tilde{\mathbb{A}}^{\mathbb{B}}[\{X\}]$ belongs to the generalized arc consistent structure \mathbb{B}' where variables in O' are fixed according to $h[O']$. Thus while enumerating such a solution $h[O']$, $ComputeAllSolutions_{DM}$ generates wrong extensions of this solution to the variable X . \triangleleft

4.2 Tight Characterizations for Enumerating Solutions with Polynomial Delay

We next analyze the complexity of $ComputeAllSolutions_{DM}$.

Theorem 5. *Let \mathbb{A} be an ℓ -structure, and $O \subseteq A$ be a set of variables. If (\mathbb{A}, O) is tp -covered through DM , then $ComputeAllSolutions_{DM}$ runs WPD.*

Proof. Assume that (\mathbb{A}, O) is tp -covered through DM . By Theorem 3, we know that $ComputeAllSolutions_{DM}$ computes the set $\mathbb{A}^{\mathbb{B}}[O]$. Thus, if the algorithm does not output any tuple, we can immediately conclude that the ECSP instance does not have solutions. Concerning the running time, we preliminary notice that the initialization steps are feasible in polynomial time. In particular, computing \mathbb{A}_V and \mathbb{B}_V is feasible in polynomial time, by the properties of the decomposition method DM (see Section 3). To characterize the complexity of the recursive invocations of *Propagate*, we have to consider instead two cases.

In the case where there is no solution, we claim that the r -structure \mathbb{B}'_V obtained by enforcing generalized arc consistency in the first invocation of *Propagate* (i.e., for $i = 1$) is empty. Indeed, since (\mathbb{A}, O) is tp -covered through DM , then $\{X_1\}$ is tp -covered in \mathbb{A}_V —just compare Definition 1 and Definition 3. It follows that we can apply Theorem 2 on the set $\{X_1\}$ in order to conclude that, for every relation $R \in \tau_V$ with $X_1 \in var(R^{\mathbb{A}_V})$, $R^{GAC(\mathbb{A}_V, \mathbb{B}_V)}[\{X_1\}] = \mathbb{A}^{\mathbb{B}}[\{X_1\}]$. Since, $\mathbb{A}^{\mathbb{B}}[O]$ is empty, the above implies that $GAC(\mathbb{A}_V, \mathbb{B}_V)$ is empty too. Thus, $ComputeAllSolutions_{DM}$ invokes just once *Propagate*, where the only operation carried out is to enforce generalized arc consistency, which is feasible in polynomial time.

Consider now the case where $\mathbb{A}^{\mathbb{B}}[O]$ is not empty. Then, the first solution is computed after m recursive calls of the procedure *Propagate*, where the dominant operation is to enforce generalized arc

consistency on the current pair $(\mathbb{A}_Y, \mathbb{B}_Y)$. In particular, by the arguments in the proof of Theorem 3, it follows that `Propagate` does not have to backtrack to find this solution: after enforcing generalized arc consistency at step i , any active value for X_i is guaranteed to occur in a solution with the current fixed values for the previous variables X_j , $1 \leq j < i$. Since GAC can be enforced in polynomial time, this solution can be computed in polynomial time as well.

To complete the proof, observe now that any solution is provided in output when `Propagate` is invoked for $i = m$. After returning a tuple of values $\langle a_1, \dots, a_m \rangle$, `Propagate` may need to backtrack to a certain index $i' \geq 1$ having some further (different) value $a_{i'}$ to be processed, fix $X_{i'}$ with this value, propagate this assignment, and continue by processing variable $X_{i'+1}$. Thus, at most m invocations of `Propagate` are needed to compute the next solution, and no backtracking step may occur before we found it. Therefore, `ComputeAllSolutionsDM` runs WPD. \triangleleft

By the above theorem and the definition of domain restricted variables, the following can easily be established.

Corollary 1. *Let \mathbf{A} be any class of ℓ -structures such that, for each $\mathbb{A} \in \mathbf{A}$, $(\mathbb{A}, \text{drv}(\mathbb{A}))$ is tp-covered through DM. Then, for every r -structure \mathbb{B} , and for every set of variables $O \subseteq \text{drv}(\mathbb{A})$, the ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ is solvable WPD.*

In the case of bounded arity structures and if the (hyper)tree width is the chosen decomposition method, it is not hard to see that the result in Corollary 1 is essentially tight. Indeed, the implication (2) \Rightarrow (1) in the theorem below easily follows from the well-known dichotomy for the decision version [17], which is obtained in the special case of ECSP instances without output variables ($O = \emptyset$).

Theorem 6. *Assume $\text{FPT} \neq \text{W}[1]$. Let \mathbf{A} be any class of ℓ -structures of bounded arity. Then, the following are equivalent:*

- (1) \mathbf{A} has bounded treewidth modulo homomorphic equivalence;
- (2) For every $\mathbb{A} \in \mathbf{A}$, for every r -structure \mathbb{B} , and for every set of variables $O \subseteq \text{drv}(\mathbb{A})$, the ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ is solvable WPD.

Actually, from an application perspective of this result, we observe that there is no efficient algorithm for the no-promise problem for such classes. In fact, the following proposition formalizes and generalizes previous observations from different authors about the impossibility of actually trusting positive answers in the (promise) decision problem [24, 4].

We say that a pair (h, c) is a certified projected solution of $(\mathbb{A}, \mathbb{B}, O)$ if, by using the certificate c , one may check in polynomial-time (w.r.t. the size of $(\mathbb{A}, \mathbb{B}, O)$) whether $h \in \mathbb{A}^{\mathbb{B}}[O]$. E.g., any full solution extending h is clearly such a certificate. If $O = \emptyset$, h is also empty, and c is intended to be a certificate that (\mathbb{A}, \mathbb{B}) is a “Yes” instance of the decision CSP. Finally, we assume that the empty output is always a certified answer, in that it entails that the input is a “No” instance, without the need for an explicit certificate of this property.

Proposition 1. *The following problem is NP-hard: Given any ECSP instance $(\mathbb{A}, \mathbb{B}, O)$, compute a certified solution in $\mathbb{A}^{\mathbb{B}}[O]$, whenever (\mathbb{A}, O) is tp-covered through DM; otherwise, there are no requirements and any output is acceptable. Hardness holds even if DM is the treewidth method with $k = 2$, the vocabulary contains just one binary relation symbol, and $O = \emptyset$.*

Proof. We show a polynomial-time Turing reduction from the NP-hard 3-colorability problem. Let M be a Turing transducer that solves the problem, that is, whenever (\mathbb{A}, O) is tp-covered through DM, at the end of a computation on a given input $(\mathbb{A}, \mathbb{B}, O)$ its output tape contains a certified solution in $\mathbb{A}^{\mathbb{B}}[O]$, otherwise, everything is acceptable. In particular, we do not pretend that M recognizes whether the above condition is fulfilled.

Then, we use M as an oracle procedure within a polynomial time algorithm that solves the 3-colorability problem. Let $G = (N, E)$ be any given graph, and assume w.l.o.g. that it contains a triangle $\{n_1, n_2\}$, $\{n_2, n_3\}$, and $\{n_3, n_1\}$. (Otherwise, select any arbitrary vertex n_1 of G and connect it to two fresh vertices n_2 and n_3 , also connected to each other. It is easy to check that this new graph is 3-colorable if, and only if, the original graph G is 3-colorable, as the two fresh vertices have no connections with the rest of the graph.) Build the (classical) binary CSP $(\mathbb{A}_G, \mathbb{B}_{3c})$ where the vocabulary contains one relation symbol R_E , and the set of variables is $A_G = N$. Moreover, $R_E^{\mathbb{A}_G} = \{\langle n_i, n_j \rangle \mid \{n_i, n_j\} \in E\}$, and $R_E^{\mathbb{B}_{3c}} = \{\langle c, c' \rangle \mid c \neq c', \{c, c'\} \subseteq \{1, 2, 3\}\}$. Consider the treewidth method for $k = 2$, and compute in polynomial time the pair $(\mathbb{A}_V, \mathbb{B}_V)$ where $\mathbb{A}_V = \ell\text{-}tw_2(\mathbb{A}_G)$ and $\mathbb{B}_V = r\text{-}tw_2(\mathbb{B}_{3c})$. In particular, observe that the hypergraph $\mathcal{H}_{\mathbb{A}_V}$ contains a hyperedge $\{n_i, n_j, n_l\}$ for every triple of vertices of G .

It is well-known and easy to see that G is 3-colorable if, and only if, $\mathbb{A}_G^{\mathbb{B}_{3c}} \neq \emptyset$, that is, if there is a homomorphism from \mathbb{A}_G to a triangle (indeed, \mathbb{B}_{3c} is a triangle). Therefore, if G is 3-colorable, the triangle substructure \mathbb{A}' such that $R_E^{\mathbb{A}'} = \{\langle n_i, n_j \rangle \mid \{i, j\} \subset \{1, 2, 3\} \mid i \neq j\}$ is homomorphically equivalent to \mathbb{A}_G . Moreover, in this case the hypergraph consisting of the single hyperedge $\{n_1, n_2, n_3\}$ is a tree projection of $\mathcal{H}_{\mathbb{A}'}$ w.r.t. $\mathcal{H}_{\mathbb{A}_V}$, or, equivalently, the treewidth of \mathbb{A}' is 2.

Now, run M on input $(\mathbb{A}_G, \mathbb{B}_{3c}, \emptyset)$ and consider its first output certificate c —say, for the sake of presentation, a full solution for the problem. Then check in polynomial time whether c is a legal certificate—in our exemplification, whether it encodes a solution of the given instance. If this is the case, we know that G is 3-colorable; otherwise, we conclude that G is not 3-colorable. Indeed, M must be correct on 3-colorable graphs, because there exists a tree projection of $\mathcal{H}_{\mathbb{A}'}$ (and thus of any core of \mathbb{A}_G —note that $O = \emptyset$ and thus there is no further requirement) w.r.t. $\mathcal{H}_{\mathbb{A}_V}$. Since all these steps are feasible in polynomial-time, we are done. \triangleleft

5 Enumerating Solutions over Arbitrary Output Variables

In this section we consider structural properties that are independent of output variables, so that tractability must hold for any desired sets of output variables. For this case, we are able to provide certified solutions WPD, which seems the more interesting notion of tractability for actual applications.

Figure 3 shows the `ComputeCertifiedSolutionsDM` algorithm computing all solutions of an ECSP instance, with a certificate for each of them. The algorithm is parametric w.r.t. any chosen decomposition method DM , and resembles in its structure the `ComputeAllSolutionsDM` algorithm. The main difference is that, after having found an assignment $\langle a_1, \dots, a_m \rangle$ for the variables in O , `ComputeCertifiedSolutionsDM` still iterates over the remaining variables in order to find a certificate for that projected solution. Of course, `ComputeCertifiedSolutionsDM` does not backtrack over the possible values to be assigned to the variables in $\{X_{m+1}, \dots, X_n\}$, since just one extension suffices to certify that this partial solution can be extended to a full one. Thus, we break the cycle after an element $\langle a_i \rangle$ is picked from its domain and correctly propagated, for each $i > m$, so that in these cases we eventually backtrack directly to $i = m$ (to look for a new projected solution).

Note that `ComputeCertifiedSolutionsDM` incrementally outputs various solutions, but it halts the computation if the current r -structure \mathbb{B}'_V becomes empty. As an important property of the algorithm, even when this abnormal exit condition occurs, we are guaranteed that all the elements provided as output until this event are indeed solutions. Moreover, if no abnormal termination occurs, then we are guaranteed that all solutions will actually be computed. Correctness follows easily from the same arguments used for `ComputeAllSolutionsDM`, by observing that, whenever $(\mathcal{H}_{\mathbb{A}}, \mathcal{H}_{\ell\text{-}DM(\mathbb{A})})$ has a tree projection, the full set of variables A is tp -covered through DM .

Theorem 7. *Let \mathbb{A} be an ℓ -structure, and $O \subseteq A$ be a set of variables. Then, for every r -structure \mathbb{B} , `ComputeCertifiedSolutionsDM` computes WPD a subset of the solutions in $\mathbb{A}^{\mathbb{B}}[O]$, with a certificate for each of them. Moreover,*

Input: An ECSP instance $(\mathbb{A}, \mathbb{B}, O)$, where $O = \{X_1, \dots, X_m\}$;
Output: for each solution $h \in \mathbb{A}^{\mathbb{B}}[O]$, a certified solution (h, h') ;
Method: let $A = \{X_1, \dots, X_m, X_{m+1}, \dots, X_n\}$ be the variables of \mathbb{A} ;
 update $(\mathbb{A}, \mathbb{B}, A)$ with any of its domain restricted versions;
 let $\mathbb{A}_V := \ell\text{-DM}(\mathbb{A})$, $\mathbb{B}_V := r\text{-DM}(\mathbb{A}, \mathbb{B})$;
 invoke $\text{CPropagate}(1, (\mathbb{A}_V, \mathbb{B}_V), m, \langle \rangle)$;

Procedure $\text{CPropagate}(i: \text{integer}, (\mathbb{A}_V, \mathbb{B}_V): \text{pair of structures}, m: \text{integer}, \langle a_1, \dots, a_{i-1} \rangle: \text{tuple of values in } A^i)$;
begin
 1. let $\mathbb{B}'_V := \text{GAC}(\mathbb{A}_V, \mathbb{B}_V)$;
 2. **if** $i > 1$ and \mathbb{B}'_V is empty **then** output “DM failure” and **HALT**;
 3. let $\text{activeValues} := \text{dom}(X_i)^{\mathbb{B}'_V}$;
 4. **for each** element $\langle a_i \rangle \in \text{activeValues}$ **do**
 5. | **if** $i = n$ **then**
 6. | | **output** the certified solution $(\langle a_1, \dots, a_m \rangle, \langle a_{m+1}, \dots, a_n \rangle)$;
 7. | **else**
 8. | | update $\text{dom}(X_i)^{\mathbb{B}'_V}$ with $\{\langle a_i \rangle\}$; /* X_i is fixed to value a_i */
 9. | | $\text{CPropagate}(i + 1, (\mathbb{A}_V, \mathbb{B}'_V), m, \langle a_1, \dots, a_{i-1}, a_i \rangle)$;
 10. | | **if** $i > m$ **then** **BREAK**;
end.

Fig. 3. Algorithm $\text{ComputeCertifiedSolutions}_{\text{DM}}$.




Fig. 4. The undirected-grid structure \mathbb{A}_u , and a mapping to one of its cores.

- If $\text{ComputeCertifiedSolutions}_{\text{DM}}$ outputs “DM failure”, then $(\mathcal{H}_{\mathbb{A}}, \mathcal{H}_{\ell\text{-DM}(\mathbb{A})})$ does not have a tree projection;
- otherwise, $\text{ComputeCertifiedSolutions}_{\text{DM}}$ computes $\text{WPD } \mathbb{A}^{\mathbb{B}}[O]$.

Moreover, we next give some evidence that, for bounded arity classes of instances, we cannot do better than this. In particular, having bounded width tree-decompositions of the full structure seems a necessary condition for the tractability of the enumeration problem WPD w.r.t. arbitrary sets of output variables (and for every r -structure).

The main gadget of the proof that tree-decompositions are necessary for tractability is based on a nice feature of grids. Figure 4 shows the basic idea for the simplest case of a relational structure \mathbb{A}_u with only one relation symbol R_u such that $R_u^{\mathbb{A}_u}$ is (the edge set of) an undirected grid. Then, any substructure \mathbb{A}_1 of \mathbb{A}_u where $R_u^{\mathbb{A}_1}$ contains just one tuple is a core of \mathbb{A}_u . However, if we consider the variant of \mathbb{A}_u where there is a domain constraint $\text{dom}(X)$ for every corner X of the grid (depicted with the circles in the figure), then the unique core is the whole structure. In fact, we next prove that this property holds for any relational structure whose Gaifman graph is a grid.

Lemma 1. *Let \mathbb{A} be an ℓ -structure whose Gaifman graph is a grid G . Moreover, let $O \subseteq A$ be the set of its four corners, and assume they are domain restricted, i.e., $O \subseteq \text{drv}(\mathbb{A})$. Then, \mathbb{A} is a core.*

Proof. Let G be such a $k_1 \times k_2$ grid, and consider any homomorphism h that maps \mathbb{A} to any of its substructures \mathbb{A}' . Since the four corners $v_{1,1}, v_{1,k_2}, v_{k_1,1}, v_{k_1,k_2}$ are domain restricted, $h(X) = X$ must hold for each of them (as $\langle X \rangle$ is the one tuple of its domain constraint $\text{dom}(X)^{\mathbb{A}}$). We say that such elements are fixed.

Consider the first row $r_1 = (v_{1,1}, v_{1,2}, \dots, v_{1,k_2})$ of G . We have seen that its endpoints, which are grid-corners, are fixed. It is easy to check that h cannot map the path r_1 to any path that is longer than r_1 . However, r_1 is the shortest path connecting the fixed endpoints $v_{1,1}$ and v_{1,k_2} , and hence it must be mapped to itself. That is, $h(X) = X$ for every element X occurring in r_1 , and thus, by the same reasoning, for every element X occurring in the last row, and in the first and the last columns of the grid. It follows that the endpoints $v_{2,1}$ and v_{2,k_2} of the second row r_2 are fixed as well, and we may apply the same argument to show that all elements occurring in r_2 are fixed, too. Eventually, row after row, we get that all elements of A are fixed, and thus the identity mapping is the only possible endomorphism for \mathbb{A} , which entails that \mathbb{A} is a core. \triangleleft

We also exploit the grid-based construction from [17], whose properties relevant to this paper may be summarized as follows.

Proposition 2 ([17]). *Let $k \geq 2$ and $K = \binom{k}{2}$, and let \mathbb{A} be any τ -structure such that the $(k \times K)$ -grid is a minor of the Gaifman graph of a core of \mathbb{A} . For any given graph G , one can compute in polynomial time (w.r.t. $\|G\|$) a τ -structure $\mathbb{B}_{\mathbb{A},G}$ such that G contains a k -clique if, and only if, there is a homomorphism from \mathbb{A} to $\mathbb{B}_{\mathbb{A},G}$.*

We can now prove the necessity of bounded treewidth for tractability WPD.

Theorem 8. *Assume $\text{FPT} \neq \text{W}[1]$. Let \mathbf{A} be any bounded-arity recursively-enumerable class of ℓ -structures closed under taking minors. Then, the following are equivalent:*

- (1) \mathbf{A} has bounded treewidth;
- (2) For every $\mathbb{A} \in \mathbf{A}$, for every r -structure \mathbb{B} , and for every set of variables $O \subseteq A$, the ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ is solvable WPD.

Proof. The fact that (1) \Rightarrow (2) holds follows by specializing Theorem 7 to the tree decomposition method. We next focus on showing that (2) \Rightarrow (1) also holds.

Let \mathbf{A} be such a bounded-arity class of ℓ -structures closed under taking minors, and having unbounded treewidth. From this latter property, by the Excluded Grid Theorem [22] it follows that every grid is a minor of the Gaifman graph of some ℓ -structure in \mathbf{A} . Moreover, because this class is closed under taking minors, every grid is actually the Gaifman graph of some ℓ -structure in \mathbf{A} .

Assume there is a deterministic Turing machine M that is able to solve with polynomial delay any ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ such that $\mathbb{A} \in \mathbf{A}$. We show that this entails the existence of an FPT algorithm to solve the $\text{W}[1]$ -hard problem $p\text{-CLIQUE}$, which of course implies $\text{FPT} = \text{W}[1]$.

Let G be a graph instance of the $p\text{-CLIQUE}$ problem, with the fixed parameter $k \geq 2$. We have to decide whether G has a clique of cardinality k . We enumerate the recursively enumerable class \mathbf{A} until we eventually find an ℓ -structure \mathbb{A} whose Gaifman graph is the $(k \times K)$ -grid. Let τ be its vocabulary. Note that searching for this structure \mathbb{A} depends on the fixed parameter k only (in particular, it is independent of G).

Let $O \subseteq A$ be the variables at the four corners of this grid, and let \mathbb{A}' be the extension of \mathbb{A} such that, for every variable $X \in O$, the vocabulary τ' of \mathbb{A}' contains the domain relation-symbol $\text{dom}(X)$. Thus, the Gaifman graph of \mathbb{A}' is the same $(k \times K)$ -grid as for \mathbb{A} , but its four corners are domain-restricted in \mathbb{A}' ($O \subseteq \text{drv}(\mathbb{A}')$). From Lemma 1, \mathbb{A}' is a core.

Recall now the grid-based construction in Proposition 2: We can build in polynomial time (w.r.t. $\|G\|$) a structure $\mathbb{B}'_{\mathbb{A}',G}$ such that there is a homomorphism from \mathbb{A}' to $\mathbb{B}'_{\mathbb{A}',G}$ if, and only if, G has a clique of cardinality k .

Consider the ECSP instance $(\mathbb{A}, \mathbb{B}, O)$ where $\mathbb{A} \in \mathbf{A}$ by construction, and \mathbb{B} is the restriction of $\mathbb{B}'_{\mathbb{A}',G}$ to the vocabulary τ . Thus, compared with $\mathbb{B}'_{\mathbb{A}',G}$, the r -structure \mathbb{B} may miss the domain constraint $\text{dom}(X)$ for some output variable $X \in O$. It is easy to see that h is a homomorphism from \mathbb{A}' to $\mathbb{B}'_{\mathbb{A}',G}$

if, and only if, h is a homomorphism from \mathbb{A} to \mathbb{B} such that, for every $X \in O$, $h(X) \in \text{dom}(X)^{\mathbb{B}'_{\mathbb{A}', G}}$. Therefore, to decide whether such a homomorphism exists (and hence to solve the clique problem), we can just enumerate WPD the set of solutions $\mathbb{A}^{\mathbb{B}}[O]$ and check whether the four domain constraints on the corners of \mathbb{A}' are satisfied by any of these solutions. Now, recall that $\mathbb{B}'_{\mathbb{A}', G}$ is built in polynomial time from G , and thus every variable may take only a polynomial number of values, and of course all combinations of four values from $\text{dom}(X)^{\mathbb{B}'_{\mathbb{A}', G}}$, $X \in O$, are polynomially many. It follows that M actually takes polynomial time for computing $\mathbb{A}^{\mathbb{B}}[O]$, and one may then check in polynomial time whether the additional domain constraints in \mathbb{A}' are satisfied or not by some solution in $\mathbb{A}^{\mathbb{B}}[O]$.

By combining the above ingredients, we got an FPT algorithm to decide whether G has a clique of cardinality k . \triangleleft

References

1. I. Adler. Tree-Related Widths of Graphs and Hypergraphs. *SIAM Journal Discrete Mathematics*, 22(1), pp. 102–123, 2008.
2. A. Atserias, A. Bulatov, and V. Dalmau. On the Power of k-Consistency, In *Proc. of ICALP'07*, pp. 279–290, 2007.
3. P.A. Bernstein and N. Goodman. The power of natural semijoins. *SIAM Journal on Computing*, 10(4), pp. 751–771, 1981.
4. A. Bulatov, V. Dalmau, M. Grohe, and D. Marx. Enumerating Homomorphism. In *Proc. of STACS'09*, pp. 231–242, 2009.
5. H. Chen and V. Dalmau. Beyond Hypertree Width: Decomposition Methods Without Decompositions. In *Proc. of CP'05*, pp. 167–181, 2005.
6. D. Cohen, P. Jeavons, and M. Gyssens. A unified theory of structural tractability for constraint satisfaction problems. *Journal of Computer and System Sciences*, 74(5): 721–743, 2008.
7. D.A. Cohen. Tractable Decision for a Constraint Language Implies Tractable Search. *Constraints*, 9(3), 219–229, 2004.
8. R. Dechter and A. Itai. Finding All Solutions if You can Find One. In *Proc. of AAAI-92 Workshop on Tractable Reasoning*, pp. 35–39, 1992.
9. R.G. Downey and M.R. Fellows. *Parameterized Complexity*. Springer, New York, 1999.
10. N. Goodman and O. Shmueli. The tree projection theorem and relational query processing. *Journal of Computer and System Sciences*, 29(3), pp. 767–786, 1984.
11. G. Gottlob, N. Leone, and F. Scarcello. A Comparison of Structural CSP Decomposition Methods. *Artificial Intelligence*, 124(2): 243–282, 2000.
12. G. Gottlob, N. Leone, and F. Scarcello. Hypertree decompositions and tractable queries. *Journal of Computer and System Sciences*, 64(3), pp. 579–627, 2002.
13. G. Gottlob, N. Leone, and F. Scarcello. Robbers, marshals, and guards: game theoretic and logical characterizations of hypertree width. *J. of Computer and System Sciences*, 66(4), pp. 775–808, 2003.
14. G. Gottlob, Z. Miklós, and T. Schwentick. Generalized hypertree decompositions: NP-hardness and tractable variants. *Journal of the ACM*, 56(6), 2009.
15. G. Greco and F. Scarcello. The Power of Tree Projections: Local Consistency, Greedy Algorithms, and Larger Islands of Tractability. To appear in *Proc. of PODS'10*.
16. M. Grohe, T. Schwentick, and L. Segoufin. When is the evaluation of conjunctive queries tractable? In *Proc. of STOC'01*, pp. 657–666, 2001.
17. M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1), 2007.
18. M. Grohe and D. Marx. Constraint solving via fractional edge covers. In *Proc. of SODA'06*, pp. 289–298, 2006.
19. D. Marx. Approximating fractional hypertree width. In *Proc. of SODA'09*, pp. 902–911, 2008.
20. D. Marx. Tractable Hypergraph Properties for Constraint Satisfaction and Conjunctive Queries. To appear in *Proc. of STOC'10*.
21. N. Robertson and P.D. Seymour. Graph minors III: Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36, pp. 49–64, 1984.

22. N. Robertson and P.D. Seymour. Graph minors V: Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41, pp. 92-114, 1986.
23. Y. Sagiv and O Shmueli. Solving Queries by Tree Projections. *ACM Transaction on Database Systems*, 18(3), pp. 487-511, 1993.
24. F. Scarcello, G. Gottlob, and G. Greco. Uniform Constraint Satisfaction Problems and Database Theory. In *Complexity of Constraints*, LNCS 5250, pp. 156-195, Springer-Verlag, 2008.